

**Amendments to the claims**

This listing of claims will replace all prior versions and listings of claims in the above-identified application.

**Listing of claims**

1. (Original) A method comprising:  
maintaining a copy of a data change log at a primary node, wherein  
said data change log at said primary node is associated with a primary data  
volume of said primary node, and  
said copy of said data change log is maintained at a data recovery node;  
detecting a failure of said primary data volume; and  
updating a secondary data volume of a secondary node using said copy of said data  
change log in response to said detecting.
2. (Original) The method of claim 1, wherein said maintaining comprises:  
maintaining a real-time copy of said data change log at said primary node.
3. (Original) The method of claim 2 further comprising:  
replicating data to be written to said primary data volume from said primary node to said  
secondary node.
4. (Original) The method of claim 3, wherein said maintaining said real-time copy  
comprises:  
receiving a request to perform a write operation on said primary data volume;  
storing data associated with said write operation substantially simultaneously on said data  
change log and said real-time copy of said data change log in response to said  
receiving.
5. (Original) The method of claim 3, wherein said replicating comprises:  
replicating said data to be written to said primary data volume to said secondary data  
volume.
6. (Original) The method of claim 3, wherein said replicating comprises:

asynchronously replicating said data to be written to said primary data volume to said secondary data volume.

7. (Original) The method of claim 3, wherein said updating comprises:  
receiving a manual update initiation indication; and  
updating said secondary data volume using said real-time copy of said data change log in response to said receiving.
8. (Original) The method of claim 3, wherein  
said real-time copy of said data change log comprises a plurality of entries; and  
said updating comprises:  
identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume, and  
updating said secondary data volume using said entry.
9. (Original) The method of claim 3, wherein said updating comprises:  
copying a block of data from said real-time copy of said data change log to a staging log at said secondary node, said block of data comprising a plurality of entries;  
applying each of said plurality of entries to a data change log at said secondary node in response to said copying; and  
updating said secondary data volume using said data change log at said secondary node.
10. (Original) The method of claim 3, further comprising:  
detecting a recovery of said primary data volume; and  
resynchronizing said primary data volume and said secondary data volume in response to said detecting.
11. (Original) A machine-readable medium having a plurality of instructions executable by a machine embodied therein, wherein said plurality of instructions when executed cause said machine to perform a method comprising:  
maintaining a copy of a data change log at a primary node, wherein said data change log at said primary node is associated with a primary data volume of said primary node, and  
said copy of said data change log is maintained at a data recovery node;

detecting a failure of said primary data volume; and  
updating a secondary data volume of a secondary node using said copy of said data  
change log in response to said detecting.

12. (Original) The machine-readable medium of claim 11, wherein said maintaining comprises:  
maintaining a real-time copy of said data change log at said primary node.
13. (Original) The machine-readable medium of claim 12, said method further comprising:  
replicating data to be written to said primary data volume from said primary node to said secondary node.
14. (Original) The machine-readable medium of claim 12, wherein said maintaining said real-time copy comprises:  
receiving a request to perform a write operation on said primary data volume;  
storing data associated with said write operation substantially simultaneously on said data change log and said real-time copy of said data change log in response to said receiving.
15. (Original) The machine-readable medium of claim 12, wherein said replicating comprises:  
asynchronously replicating said data to be written to said primary data volume to said secondary data volume.
16. (Original) The machine-readable medium of claim 12, wherein said updating comprises:  
receiving a manual update initiation indication; and  
updating said secondary data volume using said real-time copy of said data change log in response to said receiving.
17. (Original) The machine-readable medium of claim 12, wherein  
said real-time copy of said data change log comprises a plurality of entries; and  
said updating comprises:  
identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume, and

updating said secondary data volume using said entry.

18. (Original) The machine-readable medium of claim 12, wherein said updating comprises:  
copying a block of data from said real-time copy of said data change log to a staging log  
at said secondary node, said block of data comprising a plurality of entries;  
applying each of said plurality of entries to a data change log at said secondary node in  
response to said copying; and  
updating said secondary data volume using said data change log at said secondary node.

19. (Original) A data processing system comprising:  
means for maintaining a copy of a data change log at a primary node, wherein  
said data change log at said primary node is associated with a primary data  
volume of said primary node, and  
said copy of said data change log is maintained at a data recovery node;  
means for detecting a failure of said primary data volume; and  
means for updating a secondary data volume of a secondary node using said copy of said  
data change log in response to a failure of said primary data volume.

20. (Original) The data processing system of claim 19, wherein said means for maintaining  
comprises:

means for maintaining a real-time copy of said data change log at said primary node.

21. (Original) The data processing system of claim 20, further comprising:  
means for replicating data to be written to said primary data volume from said primary  
node to said secondary node.

22. (Original) The data processing system of claim 21, wherein said means for maintaining  
said real-time copy comprises:

means for storing data associated with a requested write operation on said primary data  
volume substantially simultaneously on said data change log and said real-time  
copy of said data change log.

23. (Original) The data processing system of claim 21, wherein said means for replicating  
comprises:

means for asynchronously replicating said data to be written to said primary data volume to said secondary data volume.

24. (Original) The data processing system of claim 21, wherein said means for updating comprises:

means for updating said secondary data volume using said real-time copy of said data change log in response to a manual update initiation indication.

25. (Original) The data processing system of claim 21, wherein said real-time copy of said data change log comprises a plurality of entries; and said means for updating comprises:

means for identifying an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume, and means for updating said secondary data volume using said entry.

26. (Original) The data processing system of claim 21, wherein said means for updating comprises:

means for copying a block of data from said real-time copy of said data change log to a staging log at said secondary node, said block of data comprising a plurality of entries;

means for applying each of said plurality of entries from said staging log to a data change log at said secondary node; and

means for updating said secondary data volume using said data change log at said secondary node.

27. (Original) A data processing system comprising:

a storage element to store a copy of a data change log at a primary node, wherein said data change log at said primary node is associated with a primary data volume of said primary node; and

a recovery module configured to update a secondary data volume of a secondary node using said copy of said data change log in response to a failure of said primary data volume.

28. The data processing system of claim 27, wherein said storage element comprises:

a storage element to store a real-time copy of said data change log at said primary node.

29. (Original) The data processing system of claim 28, further comprising:  
a volume management module configured to mirror data to be written to said data change log to said real-time copy of said data change log.
30. (Original) The data processing system of claim 28, further comprising:  
a volume replication module configured to synchronously replicate data to be written to said data change log to said real-time copy of said data change log.
31. (Original) The data processing system of claim 28, wherein  
said real-time copy of said data change log comprises a plurality of entries; and  
said recovery module comprises a failover management module configured to identify an entry of said plurality of entries as corresponding to an incomplete write operation on said primary data volume and update said secondary data volume using said entry.
32. (Original) The data processing system of claim 28, wherein said recovery module comprises:  
a failover management module configured to copy a block of data comprising a plurality of entries from said real-time copy of said data change log to a staging log at said secondary node.
33. (Original) The data processing system of claim 28, wherein said primary node comprises a volume replication module configured to replicate data to be written to said primary data volume to said secondary node.
34. (Original) The data processing system of claim 33, wherein said volume replication module comprises a volume replication module configured to asynchronously replicate data to be written to said primary data volume to said secondary data volume.
35. (Original) A method comprising:  
maintaining a copy of a data change log at a primary node, wherein  
said data change log at said primary node is associated with a primary data volume of said primary node, and

said copy of said data change log is maintained at a data recovery node;; and replicating data to be written to said primary data volume from said primary node to said secondary node.

36. (Original) The method of claim 35, wherein said maintaining comprises: maintaining a real-time copy of said data change log at said primary node.
37. (Original) The method of claim 36, further comprising: detecting a failure of said primary data volume; and updating a secondary data volume of said secondary node using said real-time copy of said data change log in response to said detecting.